

Semantic Visualization Of Oncology Knowledge Sources

William G. Cole PhD, Lexical Technology Inc./Information Design Seattle

David D. Sherertz, Mark S. Tuttle, Goang-Tay Hsu, Lexical Technology, Inc.

Larry M. Fagan, MD, PhD, Robert W. Carlson, MD, Stanford University

Visualization of knowledge sources can have a substantial impact on the use of such sources at the point of care. This is because barriers to use at the point of care include hours required to master the electronic interfaces to those sources, and minutes required to master the electronic interfaces to those sources, and minutes required to accomplish any one retrieval. For a system to be used regularly at the point of care, therefore, it must be intuitive and fast. This paper presents a three dimensional interface to oncology knowledge sources that aims to meet this challenge.

INTRODUCTION

Information overload is a growing problem in medical care today. The past few years have seen substantial growth in electronic knowledge resources to help caregivers maintain and even improve quality of care. Medical textbooks can now be consulted from bedside

terminals. The current medical literature on a disease can be examined from a notebook computer. It is even possible to see multimedia presentations — videos or sounds — explaining symptoms or procedures better than any written document ever could. Because these new resources are electronic, they can in principle be brought directly to the point of care: the bedside, a nursing station, a physician's private office. Every day, questions arise that could be answered by these electronic knowledge resources. "What is a tamoxifen flare? How do you stage a paranasal sinus cancer? Are there any protocols for the cancer my patient has and is there a local site my patient can go to?" But much of this medical knowledge is never consulted at the point of care. The questions may get answered, but by more traditional means — perhaps by a trip to the medical library, perhaps by telephoning a colleague. Why aren't electronic knowledge resources used more?

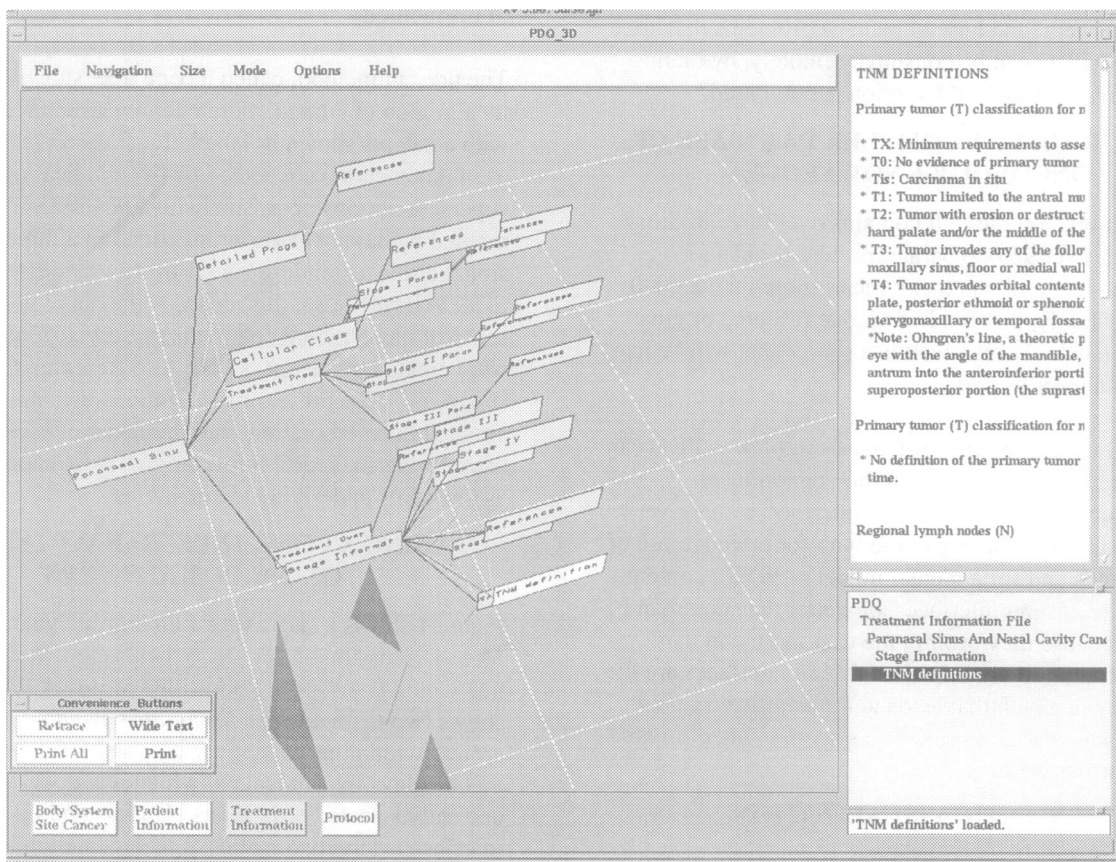


Figure 1. A semantic visualization called a cone tree.

There are two key roadblocks. First, it takes a certain amount of cognitive effort to use these electronic knowledge resources. Second, getting to this knowledge takes time. But doctors don't want to become computer experts and they don't have the spare time to slowly work their way through electronic knowledge resources, especially not at the point of care.

There are promising developments on the horizon, though. New computer hardware and software is emerging that allows caregivers to quickly and intuitively get answers from electronic knowledge resources. The idea is that a caregiver is much more likely to use a system if it quickly leads to an answer and its use is intuitively obvious. Quick — because the point of care is time limited and action oriented. Intuitive — because a doctor is not interested in learning a manual full of commands and procedures.

Figure 1 shows one new approach to presenting knowledge at the point of care: a *cone tree* interface to an oncology knowledge resource. Created by the National Cancer Institute, this valuable electronic knowledge resource called *PDQ* is a constantly updated storehouse of state of the art knowledge on how to diagnose cancers, how to treat them, and how to manage problems that arise during care. *PDQ* can be consulted remotely, by dialing in, or locally, if a CD-ROM version is purchased.

PDQ: A RICH BUT COMPLEX DATABASE OF ONCOLOGY KNOWLEDGE

PDQ is rich but complex. It contains over three million lines of text, over 95,000 paragraph-like chunks of knowledge. It is a hierarchical tree of more than 25,000 branches. Clearly, this is a knowledge resource of great promise, but equally clearly, it is a source that will be hard for the casual user to grasp.

Getting full use out of any system as complex as *PDQ* can be difficult. Using a lowest common denominator approach such as a command line interface, a caregiver would find a steep learning curve both for effective use of the commands and for a mental model of *PDQ*'s content and structure. Menu interfaces are better than command line ones because they prompt the user more and they present a structured organization. Menu interfaces are the first step down a path that leads to a *semantic visualization* of the knowledge space. The system being presented here is a further step.

SEMANTIC VISUALIZATION.

Semantic visualization is explicitly contrasted with the more commonly encountered term *scientific visualization*,

which typically refers to computer reconstruction of real world objects or events from data collected on those objects or events. A prototypical scientific visualization is an animated computer reconstruction of a powerful thunderstorm over the plains of Kansas. Semantic visualization, on the other hand, is visualization of a semantic or conceptual domain rather than an object. If we used a computer to construct a picture of a clinical database, for example, it would not look like a file cabinet or a computer disk, or any other physical object. The greatest challenge of semantic visualization is often, "Well, what does X look like?", where X is a clinical database, a body of literature, or a structured database of related concepts such as *PDQ*.

Previous authors have suggested that a domain of interrelated concepts might look like a net, with nodes or cards representing individual concepts and the lines of the net indicating relationships^[1]. Others have suggested the concept of a fisheye view^[2], a visualization in which the area of the net around whichever item is in focus is shown in great detail but areas of the net progressively farther away are depicted in progressively less detail. Greenes^[3] proposed just such a visualization for medical concepts, and the present visualization is in many ways an extension of his method.

The present method, called a cone tree, shows in detail a local region of a large, hierarchically structured network, with concepts shown as labelled rectangles (called *cards*, as in index cards) connected by lines indicating subset/superset relationships. The fisheye feature is implemented by having cards wrap around in a three dimensional circle, extending backwards in the depth plane. As in all fisheye views, this means some detail is lost, but a sense of navigational context is preserved. Cone trees were created by Xerox 's Palo Alto Research Center (PARC) ^[4] The present visualization system was constructed in collaboration with members of PARC, and uses the Xerox PARC Information Visualization Toolkit as one of its underlying technologies.

USING A CONE TREE TO NAVIGATE TO ONCOLOGY ANSWERS

Figure 1 shows a static view of a highly dynamic interface. Discussion of this figure helps explain how a cone tree would be used for knowledge retrieval. Figure 1 shows two main areas: a *Navigation Window* for moving around within the semantic space, and a *Content Window* that displays the contents of whichever node or card is currently selected or in focus. Understanding the entire interface is largely a matter of understanding how the *Navigation Window* causes changes in a second window, called the *Content Window*.

Clicking on labelled cards in the Navigation Window causes the system to move up, down, or laterally in the PDQ hierarchy. At every moment one and only one card is the focus of attention. This card is highlighted on the screen, to signal to the user which card is the focus.

The *Content Window* is displayed immediately to the right of the Navigation Window. This window displays whatever content PDQ has available for the topic that is the current focus. If PDQ has no content available specific to this topic (e.g. if the topic is very high level, such as “Head And Neck Cancer”), the Content Window simply displays “No content available for this topic.” In Figure 1, the caregiver has navigated from the highest level of cancers (where the names of the major cancer divisions — Gastrointestinal, Head And Neck, and so on — are displayed) down to a subset of the Paranasal Sinus area of PDQ’s Information For Physicians (sometimes called “State Of The Art” information). The topic of focus is “TNM definitions”, the topic in which the TNM

(Tumor, Nodes, Metastases) method of staging a cancer is explained. The Content Window displays PDQ’s content on this topic in a scrollable list.

A few frequently used features are accessible by buttons prominently displayed on the lower left. Figure 2 shows what happens when one of these buttons, the Wide Text button, is clicked. Until the button is clicked again, content is favored over navigation, making it much easier to read the content text.

PRINTING INFORMATION

After the oncologist has used this system to find information tailored for physicians, pressing another of the frequently used buttons, Patient Information, carries the oncologist over to a parallel tree, where information on the topic of focus is found, but tailored for presentation to patients rather than physicians. Access to this type of information is one of the most useful parts of PDQ, as it helps both the physician and the patient, by providing the patient with an understandable account of what is known about his or her cancer, how it will be treated and why.

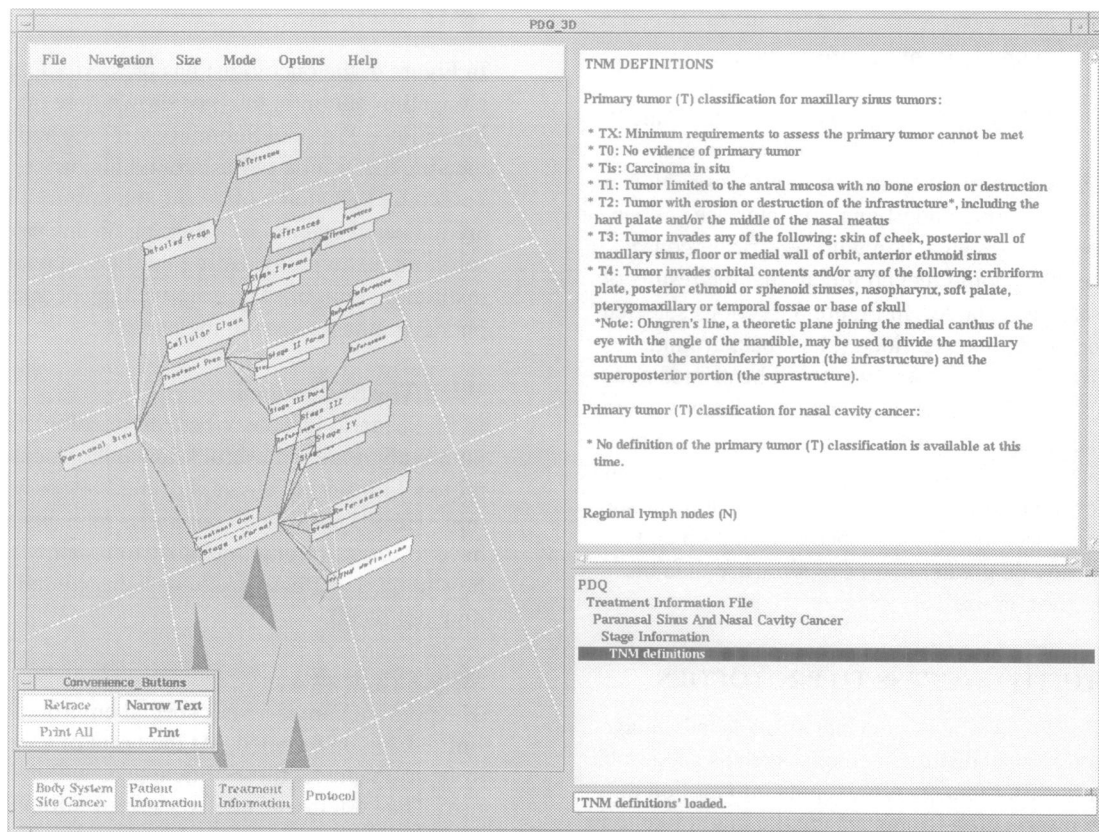


Figure 2. The Wide Text view favors content over navigation.

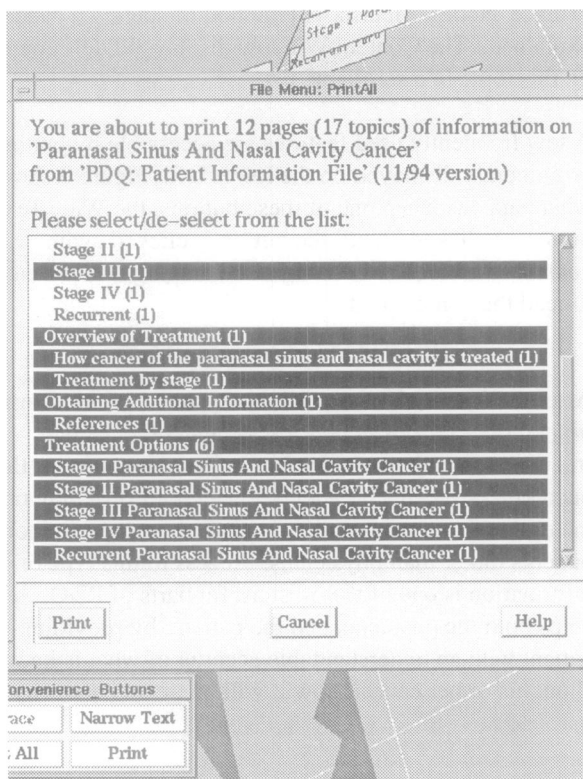


Figure 3. Printing information.

In Figure 3 the oncologist has moved over to the Patient Information tree by pressing the Patient Information button and is now putting a print job in progress. Note a key cognitive principle underlying this interface. When the user contemplates an action with significant consequences, the user should feel confident that the results of this action are predictable. In Figure 3, sections selected to be printed are in reverse video (black background, white lettering). The user sees what will be printed and also sees a summary statement saying how many sections and how many pages this print job represents. Accomplishing this action can be quick only if it is intuitive, and that requires the user feel confident about what he or she is doing. Feedback is crucial to speed; hesitation and uncertainty may be greater causes of slow use than the speed of the system itself.

INTUITIVE ACCESS TO PROTOCOLS

A semantic visualization system can also take advantage of more typical visualization methods, such as geographic visualization. After this oncologist puts into motion a print job of information for his patient, he might like to seek out an experimental treatment protocol for his patient. Since the structure of PDQ does not store protocol information as part of a cancer description, the

oncologist would need to move to a separate tree, as he did when going to Patient Information. And, as with Patient Information, this would be accomplished via one of the frequently used buttons visible in Figure 1, the one marked *Protocol*.

Clicking on this button would take the physician to a Protocol Summary window where information about protocols appropriate for the cancer under consideration would be found. Two distinct types of action can be taken in this section. The first, not illustrated here, is navigating through the many requirements, restrictions, and descriptions of the protocol. This action uses the standard PDQ-3D interface shown in Figures 1 and 2 — a Navigation Window with cone tree for moving around, a Content Window to display text contained at nodes.

Once a protocol has been found that looks to be a likely clinical match for this patient, the oncologist will need to ask some logistical questions. Is there a nearby site that offers this protocol? Who is in charge of the protocol at that site? Are they still enrolling patients and if so, what is the phone number I should call? Support for answering these is available through a geographic navigation display, the Protocol Map.

In Figure 4, the oncologist has pressed a *Find Local* navigation button in the (not shown here due to space limitations) Protocol Summary window and is now viewing a 300 mile radius around his office. Institutions known by PDQ to be offering the protocol in question are indicated by circles. Clicking on any one brings up a scrolling text window listing the PDQ physicians at that institution, with brief descriptive tags of their roles, and phone numbers to enroll the patient.

What The System Did

In the scenario just described, an oncologist quickly and intuitively retrieved useful knowledge, located a nearby protocol for his patient, and printed out patient information. He did not consult computer manuals nor did he have to search through memory to recall what to do or how to do it. Effective actions were guided by recognition, not recall.

Why This Matters

The point of care is a very special environment, time limited, action oriented. At the point of care, only systems that meet three critical tests will ever see frequent use. The system must deliver useful information; it must do so quickly; and it must place low demands on human information processing. The system described here aims to meet these tests.

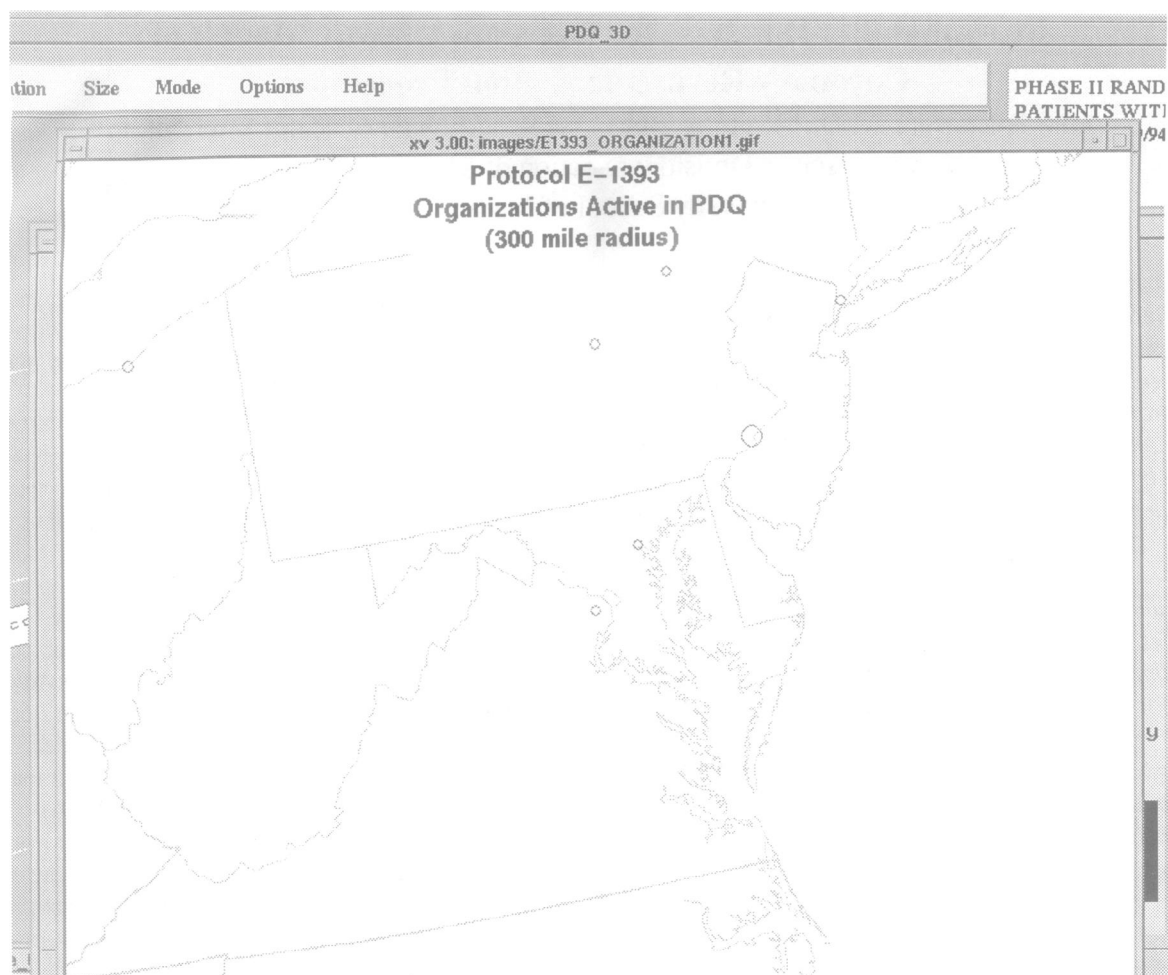


Figure 4. The Protocol Map. Protocol location through geographic visualization.

DEVELOPMENT AND EVALUATION PLAN

This paper reports on Phase I of a three phase project sponsored by the National Cancer Institute. In Phase II, the prototype system briefly sketched here will be extended in several ways and evaluated. Content will be expanded to cover CANCERLIT and the textbook *Principles and Practices of Oncology (P&PO)*. Improvements in the cone tree representation will focus on refining cues to structure and function such as maximizing the visual cues supporting the three dimensional illusion. Implementation of a second 3D semantic visualization will support linear structures such as a database of journal articles and other knowledge in which time is a dimension of primary concern. The evaluation plan proposed is to have oncology fellows learn the system then carry out typical knowledge retrieval tasks under laboratory conditions. We will measure learning time, retention of learning over time, understanding of the structure of the knowledge source, navigational ability,

and similar cognitive components of knowledge retrieval. Phase II will begin in October 1995 and run through October 1997.

References

- [1]. Fairchild KM, Poltrock SE, and Furnas GW. SemNet: Three dimensional graphic representation of large knowledge bases. Appearing in: *Cognitive science and its applications for human-computer interactions*, R. Guindon, Ed. Newark: Erlbaum, 1988.
- [2]. Furnas, G. Generalized fisheye views. *Proc of CHI '86*, Boston: ACM Press, M. Mantei Ed., 1986.
- [3]. Komorowski HJ, Greenes RA, and Pattison-Gordon E. The use of fisheye views for displaying semantic relationships in a medical taxonomy. *Proc of SCAMC*, Washington: ACM Press, W. Stead Ed., 1987.
- [4]. Robertson GG, Card SS, and Mackinlay JD. Information visualization using 3D interactive animation. *Comm of ACM*, 36(4), 57-71, 1993.